Google Inc.
1101 New York Avenue, N.W.
Second Floor
Washington, DC 20005

Main 202 346.1100
Fax 202 346.1101
www.google.com

May 9, 2011

The Honorable Fred Upton
Chairman, Committee on Energy and Commerce

The Honorable Lee Terry
Vice Chairman, Subcommittee on Communications and Technology

The Honorable Greg Walden
Chairman, Subcommittee on Communications and Technology

The Honorable Mary Bono Mack
Chair, Subcommittee on Commerce, Manufacturing and Trade

The Honorable Marsha Blackburn
Vice Chair, Subcommittee on Commerce, Manufacturing and Trade

United States House of Representatives
Committee on Energy and Commerce
Washington, DC 20515

Dear Chairman Upton, Vice Chairman Terry, Chairman Walden, Chair Bono Mack, and Vice Chair Blackburn:

This letter responds to your letter dated April 25, 2011 to Google Inc. (addressed to Larry Page, CEO), which requests information related to the collection and storage of location information by devices running the Android operating system. Google appreciates the opportunity to respond to your questions, and we hope to correct some of the inaccurate information reported in the media about Android and location information. Below we have provided an overview of our location practices for Android, and then we address the specific questions from your letter.

Introduction

Location-based services are some of the most exciting and innovative new products in the Internet marketplace today. These services are already providing consumers with enormous benefits and are sure to become more useful in the future. From navigation tools (such as maps) to localized search to social applications, consumers and businesses are increasingly benefiting from Internet services that are enhanced by location.

Many companies have created products that rely on a user's estimated location to provide a more useful product experience. Twitter allows users to "geotag" their tweets, because knowing someone is tweeting from Bangkok as opposed to Manchester gives followers important context and perspective. On smartphones like iPhone, Palm, and Android devices, services such as Yelp and Urbanspoon use location to provide relevant local search results, while applications like Foursquare let users find nearby friends.

At Google, we have seen an explosion in demand for location-based services. For example, users can use our location-based services to find driving directions from their current location, identify a traffic jam and find an alternate route, and find the next movie time at a nearby theater. In the last year, 40% of Google Maps usage was from mobile devices. There are now 150 million active monthly Google Maps for mobile users on Android, iPhone, BlackBerry, and other mobile platforms in more than 100 countries. None of this would have been possible without the use of location information.

While location-based services are already showing great value to users, Google recognizes the particular privacy concerns that come with the collection and storage of location information. That is why we don't collect any location information through our location provider on Android devices unless the user specifically "opts-in" to share this information with Google. Even after opting-in, we give users a way to easily turn off location sharing with Google at any time they wish.

Google is also very careful about how we use and store the data that is generated by location-based services. Our users contribute to location-based services through Android devices by opting-in to the collection of their location information. We provide users with notice and control over the collection, sharing, and use of location information in order to provide a better mobile experience on Android devices. The location information sent to Google servers is anonymized and stored in the aggregate and is not tied or traceable to a specific user. A small amount of location information regarding nearby Wi-Fi access points and cell towers is kept on the device to help the user continue to enjoy the service when no server connection is available and to improve speed and battery life. This information on the device is likewise not tied or traceable to a specific user.

In order to give some context to the answers to your specific questions, we think that it will be helpful to explain how Android provides location-based services. So below we have described the design of the Android operating system, the Google Location Server (our mapping of Wi-Fi access points and cell towers), the Google Network Location Provider application for Android, and how users can choose to share location information with third party applications on Android.

Android Operating System Design

Android is an open source operating system designed for mobile devices. The open source nature of the Android operating system brings equipment manufacturers, carriers, application developers, and end users unprecedented freedom of choice at all levels of the mobile device distribution chain. The first Android device was released to the public less than three years ago, and already there have been over 170 compatible Android device models released from a network of 27 OEMs throughout the world.

Because of the open source nature of the Android operating system, device manufacturers can enable Android devices to obtain location information from a variety of sources (referred to as "location providers"). For example, most devices are equipped with a GPS receiver and that is one source of location information. The Google Location Server ("GLS") is another possible source, as implemented through the Google Network Location Provider application for Android ("NLP") which is described below. There are also third party sources of location information like Skyhook.

Location Manager is a system component of the open source Android operating system that manages the interactions between each "location provider" enabled for the device by the device manufacturer. Each location provider that is turned on by the user will periodically send a location fix to the Location Manager, consisting of latitude, longitude, and other related information such as accuracy (e.g. 1000 meter error radius). The Location Manager will temporarily store on the device the most recent location fix and will discard the previous location fix as soon as it receives an updated location fix from the particular location provider. Android applications can access the most recent location fix from the Location Manager if that application has the necessary user permission as described below in the "Android Application Permissions" section.

Because of the open source nature of the Android operating system, a device manufacturer can build an Android device without any involvement by Google. Device manufacturers may also choose to enable non-Google location providers (e.g. Skyhook). Google has no control over how non-Google location providers are implemented on devices. Therefore, the responses in this letter refer only to unmodified versions of the Android operating system as released by Google and the GLS and NLP location provider created by Google.

GLS - Generally

Location providers can rely on a variety of location indicators associated with the user's device to help estimate the device's location. For example, GPS enabled devices can provide a highly accurate location using information from GPS satellites. But GPS can be slow and drain battery life and can take tens of seconds and sometimes much longer to "fix" a location depending on the specific hardware and the physical location. Furthermore, many devices are not GPS enabled or are used in situations where obtaining a GPS signal might not even be possible (e.g. indoors, where there is no line of sight between the device and the satellites).

As a result, various companies have worked out other solutions as an alternative to GPS. These are generally based around the idea of detecting nearby, publicly available signals from Wi-Fi access points and cell towers and using this data to quickly approximate a rough position, usually with less accuracy than GPS. By treating Wi-Fi access points or cell towers as beacons, devices are able to fix their general location quickly in a power-efficient way, even while they may be working on a more precise GPS-based location. This can be done by using information that is publicly broadcast (for example, that list of Wi-Fi access points you see when you use the "join network" option on your computer). A database of known network locations is required to determine a user's estimated location from either Wi-Fi access point or cell tower information. Companies like Skyhook Wireless and Navizon compile such databases and license the data to many industry leaders.

Google has also created such a location provider known as the Google Location Server or GLS. GLS is an Internet database on Google servers that uses Wi-Fi access points and cell towers to determine an estimated location and that uses GPS information to estimate road traffic.

GLS - NLP for Android

The Google Network Location Provider application for Android (or NLP) is a proprietary Google application that may be installed on Android devices by a device manufacturer pursuant to a license with Google. NLP interacts with GLS to determine a user's estimated location using Wi-Fi access points and cell towers, providing location information in a way that works indoors and outdoors, responds faster, and uses less battery power than GPS services.

Opt-in and Opt-out for NLP

NLP is off by default. The user can opt-in to turn on NLP either through the initial setup flow or from within the device settings. If a user chooses to opt-in, the user can later opt-out and turn off NLP at any time within the device settings. The following screenshot sets demonstrate the opt-in and opt-out process.

Screenshot Set #1 - Initial Setup Flow (Opt-In)

NLP is off by default. The user can opt-in to turn on NLP through the initial setup flow.[1]

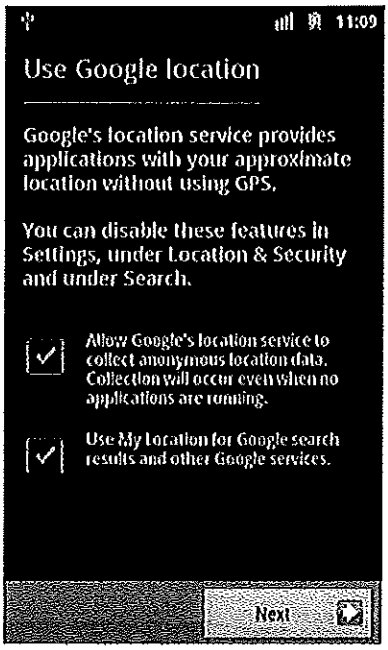

[1] Note that some device manufacturers have chosen not to present the initial setup flow on some devices (such as the Sony Ericsson Xperia X10 and LG Optimus S) in which case NLP remains off until the user turns it on by opting-in through the device settings as shown in Screenshot Set #2.

Screenshot Set #2 - Settings (Opt-In)

NLP is off by default. The user can opt-in to turn on NLP from within the device settings. The first screen shows that NLP is turned off, because there is no checkmark next to "Use wireless networks." To opt-in, the user touches the screen in the "Use wireless networks" section which causes the "Location consent" pop-up message to be displayed as shown in the second screen. If the user touches "Disagree" then the device reverts to the first screen and NLP remains off. If the user touches "Agree" then NLP is turned on, and the third screen reflects that NLP is turned on by showing the checkmark next to "Use wireless networks."

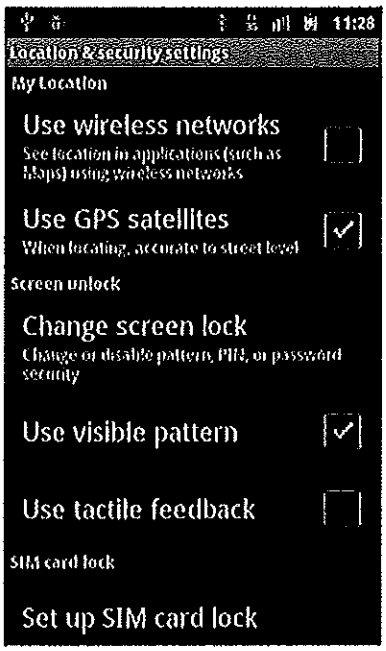

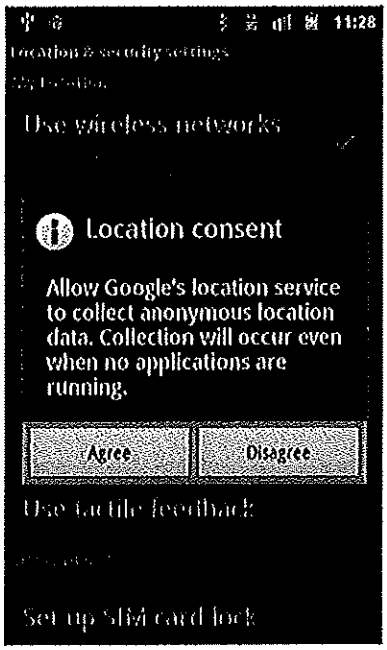

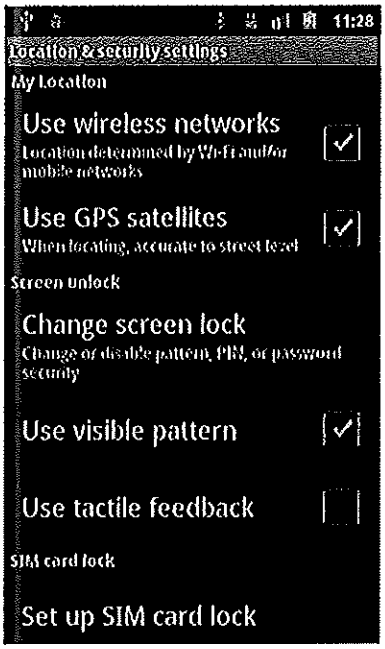

Screenshot Set #3 - Settings (Subsequent Opt-Out)

Because NLP is off by default, the opting-out process is only relevant if the user has previously opted-in to use NLP. The first screen shows that NLP is turned on, because there is a checkmark next to "Use wireless networks." To opt-out, the user simply touches the screen in the "Use wireless networks" section and the checkmark is removed. The second screen shows that NLP is now turned off.

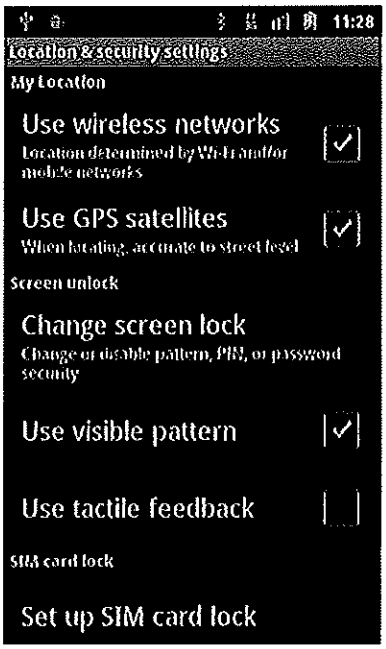

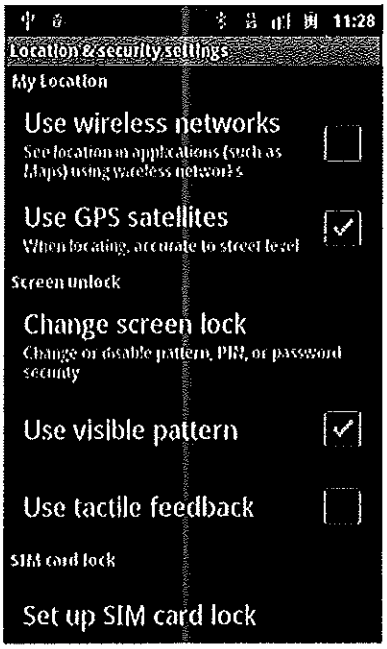

Information Transmitted to GLS

If the user has opted-in, NLP will periodically transmit the following information to GLS:

- Identification of nearby Wi-Fi access points (including SSID, BSSID, signal strength).
- Identification of nearby cell towers (including signal strength).
- GPS location and velocity, if GPS is also turned on. GPS information is associated with the identified Wi-Fi access points and cell towers and used to help estimate the location of the Wi-Fi access points and cell towers. GPS information is also used to estimate road traffic.
- Timestamp of the transmission.
- Anonymous token: Although this token is persistent and unique to the device, the token is randomly generated by GLS and has no association with any other information that could identify the particular device or particular user. A new token will be generated by GLS if the user performs a factory reset of the device. This token is only used to tag communications between NLP and GLS, which enables Google to compute velocity for road traffic estimates and to identify invalid transmissions to GLS. Google wants only legitimate devices to provide information, because invalid transmissions have the potential to pollute the GLS database with inaccurate data and degrade the ability of GLS to provide reliable location information to future users.
- IP address of the device: For devices operating on a cellular data network, this IP address is typically a dynamically assigned gateway or proxy IP address from the mobile carrier, which is shared by many users. IP address is also used to identify invalid transmissions to GLS (for the same reasons described in the bullet above).
- Other general information such as the device model, device wireless radio type (e.g. 802.11 a/b/g/n), user agent (e.g. Android operating system version), wireless carrier, carrier network type, battery level, status of the NLP Cache (as defined below), the application that is invoking the Location Manager to make the request to NLP, and information about the request being made by the application.

The collected information is stored in temporary databases on Google servers for approximately one week in association with a hashed version of the anonymous token. The token is put through a one-way hash as soon as it arrives at the server. The token itself is never stored on a Google server, only a hash of the token is temporarily stored in which the hash key is rotated at least every seven days. After this approximately one week of temporary storage, the information related to the hashed token values and the IP addresses are stripped from the data and measures are taken to obfuscate GPS route endpoints. The remaining data is transferred into aggregate and anonymous databases on Google servers for permanent storage, consisting of a database of Wi-Fi access point and cell tower locations and a database of road traffic information.

When NLP on a user's device makes a request to the GLS server with a list of Wi-Fi access points and cell towers that are visible to the device, GLS compares these Wi-Fi access points and cell towers with the list of known Wi-Fi access points and cell towers in the GLS database. If GLS has location information for these Wi-Fi access points and cell towers, GLS uses this information to determine the approximate location of the user and returns this location fix to NLP. Simultaneously, GLS sends the location information (i.e.

latitude/longitude) of the reported Wi-Fi access points and cell towers back to the device, so that this information can be stored in a cache file (the "NLP Cache") on the device.

Information Cached on the Device

The NLP Cache contains the estimated location (latitude and longitude) for the device's most recently seen Wi-Fi access points and cell towers, along with a timestamp. For clarity, the information stored in the NLP Cache is not the user's actual location, but rather the estimated locations of the Wi-Fi access points and cell towers that were near the user and seen by the device.

The NLP Cache is designed to hold information on a maximum of 200 nearby Wi-Fi access points (or 400 on some versions of NLP) and on a maximum of 50 nearby cell towers. Because this NLP Cache has a limited capacity, older data is dropped as newer data is added. For some context to these limits on the size of the NLP Cache, a user could feasibly encounter hundreds of access points during a typical commute in an urban area. The purpose of the NLP Cache is to enable NLP to compute a location when no server connection is available and to improve speed and battery life by avoiding unnecessary server communication with GLS.

Although the NLP Cache is not encrypted, it is protected by the security architecture of the Android operating system. The NLP Cache cannot be accessed by the user and cannot be accessed by any applications on the device, unless the user has chosen to "root" the device (which disables the security architecture and gives the user full control and access to the device). The NLP Cache resides only on the device, and Google does not provide any mechanism to backup or sync the NLP Cache to a server, personal computer, or any other device.

When a user opts-out of NLP, no new data is added to the NLP Cache. In recent versions of NLP (released in late 2010) that are used with Android 2.3 and later, existing data in the NLP Cache is automatically deleted if the user later opts-out of location sharing with Google by turning off NLP. In older versions of NLP and for older versions of the Android operating system, the user could clear the NLP Cache after opting-out by performing a factory reset of the device.

Android Application Permissions

The Android operating system is built on the principle of openness, with the goal of encouraging innovation and user choice. With this principle in mind, Google does not decide which applications can access location or other user information from the device. Instead, the Android operating system uses a permissions model in which the user is automatically informed of certain types of information an application will be able to access during the application installation process. This permissions model is designed to empower users to make their own decision on whether or not to trust an application with the information requested. The user may choose to trust the application by completing the installation or the user may choose to cancel the installation.

An application can access the device's GPS location through a permission, which will display the permission message "Your location: fine (GPS) location" to the user during

installation. An application can access the device's network location (e.g. approximate location from NLP and GLS) through a permission, which will display the permission message "Your location: coarse (network-based) location" to the user during installation. An application can obtain Internet access through a permission, which will display the message "Network communication: full Internet access" to the user during installation.

When Google is the developer of the Android application (e.g. the Google Maps for mobile application), Google is responsible for how the application collects and handles data and for the privacy disclosures made to users. Most Google-developed Android applications are subject to the Google Mobile Terms of Service at http://m.google.com/tos and the Google Mobile Privacy Policy at http://wwww.google.com/intl/en/mobile/privacy.html, unless Google has created a custom terms of service and privacy policy for the application. Google privacy policies are also clearly displayed to the user when the user first signs into the Android device.

When Google is not the developer, the application developer bears the responsibility for the design of the application, which includes responsibility for how the application collects and handles user data and the privacy disclosures made to users. If the user chooses to trust an application with location information by proceeding with the installation after viewing the location-related permissions, then that application could potentially store this location information on the device or transmit the information off the device if the application also has the Internet access permission. Google does not control the behavior of third party applications or how they handle location information and other user information that the third party application obtains from the device. Even though the developer bears the responsibility, Google strongly encourages application developers to use best practices as described in this Google blog post: http://android-developers.blogspot.com/2010/08/best-practices-for-handling-android.html.

We hope this overview of our location practices for Android, together with our attached answers to your specific questions, addresses your concerns. However, if you would like to discuss this further, please let us know.

Kind regards,

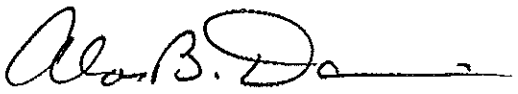

Alan B. Davidson
Director, Public Policy
Google Inc.

Attachment: Specific Questions

Below is Google Inc.'s response to the request for information from your letter dated April 25, 2011 to Google Inc. These answers should be read together with the overview in our attached cover letter.

1. What location data do devices running your operating system track, use, store, or share?

Please see the "Android Operating System Design" section above for a description of how location information is managed by the Android operating system. Please see the "GLS - NLP for Android" section above for a description of how certain location information is transmitted between NLP and GLS and how that information is used. Please see the "Android Application Permissions" section above for how users can choose to share location information with third party applications on Android by granting certain permissions.

2. Why does the device track, use, store, or share that data?

Please see the entirety of the overview of our location practices for Android devices in the attached cover letter.

3. Where on the device is the data stored; how is it used, stored, or shared; and how is it protected?

Please see the "Android Operating System Design" section above. The Location Manager will temporarily store on the device the most recent location fix and will discard the previous location fix as soon as it receives an updated location fix from the particular location provider. Android applications can access the most recent location fix from the Location Manager if that application has the necessary user permission as described in the "Android Application Permissions" section.

Please see the "GLS - NLP for Android" section above. The NLP Cache contains the estimated location (latitude and longitude) for the device's most recently seen Wi-Fi access points and cell towers, along with a timestamp. For clarity, the information stored in the NLP Cache is not the user's actual location, but rather the estimated locations of the Wi-Fi access points and cell towers that were near the user and seen by the device. The NLP Cache is designed to hold information on a maximum of 200 nearby Wi-Fi access points (or 400 on some versions of NLP) and on a maximum of 50 nearby cell towers. Because this NLP Cache has a limited capacity, older data is dropped as newer data is added. The purpose of the NLP Cache is to enable NLP to compute a location when no server connection is available and to improve speed and battery life by avoiding unnecessary server communication with GLS. Although the NLP Cache is not encrypted, it is protected by the security architecture of the Android operating system. The NLP Cache cannot be accessed by the user and cannot be accessed by any applications on the device, unless the user has chosen to "root" the device which disables the security architecture and gives the user full control and access to the device. The NLP Cache resides only on the device, and Google does not provide any mechanism to backup or sync the NLP Cache to a server, personal computer, or any other device.

4. How is that data accessible and who can access it? Is the data automatically transferred to your company or to other devices, or to third parties? If so, how and why? Is there any other manner in which the data can be transferred to or obtained by your company, or by other devices, or by third parties and, if so, how and why?

Please see answer to Question 3 above.

Also, please see the "GLS - NLP for Android" section above for a description of how certain location information is transmitted between NLP and GLS and how that information is used. NLP does not transmit this information to any third party. The Location Manager will temporarily store on the device the most recent location fix and will discard the previous location fix as soon as it receives an updated location fix from the particular location provider (such as NLP). Android applications (i.e. a third party) can access the most recent location fix from the Location Manager if that application has the necessary user permission as described in the "Android Application Permissions" section.

Device manufacturers may also choose to enable non-Google location providers (e.g. Skyhook). Google has no control over how non-Google location providers are implemented on devices or to whom these non-Google location providers transmit information.

5. Is the user informed of, or given an opportunity to prevent, such tracking, use, storing, or sharing of data and, if so, how? Can the end-user disable the tracking, use, storing, and sharing of such data? Can the user delete the data?

NLP is off by default. The user can opt-in to turn on NLP either through the initial setup flow or from within the device settings. If a user chooses to opt-in, the user can later opt-out and turn off NLP at any time within the device settings. The screenshot sets shown above demonstrate the opt-in and opt-out process. Note that some device manufacturers have chosen not to present the initial setup flow on some devices (such as the Sony Ericsson Xperia X10 and LG Optimus S) in which case NLP remains off until the user turns it on by opting-in through the device settings.

When a user opts-out of NLP, no new data is added to the NLP Cache. In recent versions of NLP (released in late 2010) that are used with Android 2.3 and later, existing data in the NLP Cache is automatically deleted if the user later opts-out of location sharing with Google by turning off NLP. In older versions of NLP and for older versions of the Android operating system, the user could clear the NLP Cache after opting-out by performing a factory reset of the device.

6. How long does the device store the data?

There is no specific time limit.

The Location Manager will temporarily store on the device the most recent location fix and will discard the previous location fix as soon as it receives an updated location fix from the particular location provider.

The NLP Cache is designed to hold information on a maximum of 200 nearby Wi-Fi access points (or 400 on some versions of NLP) and on a maximum of 50 nearby cell

towers. Because this NLP Cache has a limited capacity, older data is dropped as newer data is added.

7. Section 222 of the Communications Act contains privacy provisions. Do those provisions apply to you? Should they? Does it make sense that similar information is afforded different privacy protections depending on what entity does the collecting and what service the data is collected from, especially since the entities collecting such information are increasingly competing against each other in today's information age?

Section 222 of Title 47 requires telecommunications carriers to protect the confidentiality of their customer's proprietary network information or "CPNI." Call location information related to a subscriber's use of a commercial mobile radio service is considered to be CPNI and protected from a carrier's access, use or disclosure without the customer's express prior authorization. Because Google is not a telecommunications carrier, these requirements do not apply to it. Nor do these requirements apply to most wireless carrier data services and applications today. For example, location information associated with the provision of information services and wireless internet access are not included within the ambit of Section 222. See Appropriate Regulatory Treatment for Broadband Access to the Internet Over Wireless Networks, Declaratory Ruling, WT Docket No. 07-53 (Rel. Mar. 23, 2007); see also CTIA - The Wireless Association, Best Practices and Guidelines for Location Based Services, available at: http://www.ctia.org/business_resources/wic/index.cfm/AID/11300. The Guidelines state that they are intended to promote and protect user privacy as new Location-Based Services ("LBS") are developed and deployed. As CTIA explains it, "Location Based Services have one thing in common regardless of the underlying technology – they rely on, use or incorporate the location of a device to provide or enhance a service. Accordingly, the Guidelines are technology-neutral and apply regardless of the technology or mobile device used or the business model employed to provide LBS (e.g., a downloaded application, a web-based service, etc.)." Google supported the development of these Guidelines. Thus, the regulatory playing field is level and there is no need to extend Section 222 to non-carriers. Moreover, Google's location services are consent-based. The user has the choice to disclose his or her location or not. Thus, the regulatory purpose of Section 222 is fulfilled today without regard to a statutory mandate. And because location information may be part of other service offerings, Google notes that the Electronic Communications Privacy Act, 18 U.S.C. 2701 et seq., provides additional protection for location information.